

COMMUNICATIONS-INSPIRED PROJECTION DESIGN WITH APPLICATION TO COMPRESSIVE SENSING

WILLIAM R. CARSON^{†§}, MINHUA CHEN[‡], MIGUEL R. D. RODRIGUES^{†¶},
ROBERT CALDERBANK[‡], AND LAWRENCE CARIN[‡]

Abstract.

We consider the recovery of an underlying signal $\mathbf{x} \in \mathbb{C}^m$ based on projection measurements of the form $\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{w}$, where $\mathbf{y} \in \mathbb{C}^\ell$ and \mathbf{w} is measurement noise; we are interested in the case $\ell \ll m$. It is assumed that the signal model $p(\mathbf{x})$ is known, and $\mathbf{w} \sim \mathcal{CN}(\mathbf{w}; \mathbf{0}, \mathbf{\Sigma}_{\mathbf{w}})$, for known $\mathbf{\Sigma}_{\mathbf{w}}$. The objective is to design a projection matrix $\mathbf{M} \in \mathbb{C}^{\ell \times m}$ to maximize key information-theoretic quantities with operational significance, including the mutual information between the signal and the projections $\mathcal{I}(\mathbf{x}; \mathbf{y})$ or the Rényi entropy of the projections $h_\alpha(\mathbf{y})$ (Shannon entropy is a special case). By capitalizing on explicit characterizations of the gradients of the information measures with respect to the projections matrix, where we also partially extend the well-known results of Palomar and Verdú from the mutual information to the Rényi entropy domain, we unveil the key operations carried out by the optimal projections designs: *mode exposure* and *mode alignment*. Experiments are considered for the case of compressive sensing (CS) applied to imagery. In this context, we provide a demonstration of the performance improvement possible through the application of the novel projection designs in relation to conventional ones, as well as justification for a fast online projections design method with which state-of-the-art adaptive CS signal recovery is achieved.

Key words. Low Resolution Imaging, Compressed Sensing, MIMO Communication, Precoder Design, Mode Alignment, Mutual Information

AMS subject classifications. 94A08, 68U10, 94A05, 94A15

1. Introduction. Compressive sensing (CS) [9, 13] has recently emerged as an important area of research in image sensing and processing. Compressive sensing has been particularly successful in multidimensional imaging applications, including magnetic resonance [29], spectral imaging [17, 45] and video [44, 22]. Conventional sensing systems typically first acquire data in an uncompressed form (*e.g.*, individual pixels in an image) and then perform compression subsequently, for storage or communication. In contrast, CS involves acquisition of the data in an already compressed form, reducing the quantity of data that need be measured in the first place. In CS the underlying signal to be measured is projected onto a set of vectors [7, 13], and one must perform an inverse problem to recover the underlying signal of interest.

There are two hallmarks of the original CS theory. First, the projection vectors were usually constituted uniformly at random. Second, the underlying signal model used to regularize the inverse problem was based on the assumption that the underlying signal could be sparsely represented in terms of an orthonormal basis or frame. However, even in some of the early CS studies, it was recognized that improved performance could be achieved with projection vectors designed to the underlying signal of interest [24, 1, 2], rather than using random projections. Further, it has recently

[†]W. R. Carson and M. R. D. Rodrigues were supported by the Fundação para a Ciência e a Tecnologia through the research project PTDC/EEA-TEL/100854/2008.

[‡]M. Chen, R. Calderbank and L. Carin are with Department of Electrical Engineering, Duke University, USA. Their work was supported in part by NSF under Grant DMS-0914892, by ONR under Grant N00014-08-1-1110, and by DARPA under grant N66001-11-1-4002 as part of the KeCom Program.

[§]W. R. Carson is now with PA Consulting Group, Cambridge Technology Centre, Melbourn, UK.

[¶]M. R. D. Rodrigues is now with Dept. E & EE, University College London, London, U.K.

A part of the material presented in this work has been published in IEEE ICASSP 2012.

been recognized that a signal model based upon sparsity is often overly primitive, and model-based CS [3], wherein improved signal models are employed, may yield improved CS performance (high-quality signal recovery with fewer projection measurements). Signal models that have been considered include the Gaussian mixture model (GMM) [11], union-of-subspace models [15], and manifold models [4].

In this paper our goal is to design CS projection matrices (“measurement kernels”) matched to a general statistical signal model. Specifically, if the underlying signal to be measured is $\mathbf{x} \in \mathbb{C}^m$, it is assumed that we have access to a general signal model, represented statistically by density function $p(\mathbf{x})$. Our objective is to design the projection matrix to maximize the mutual information between the underlying signal and the observed compressive measurements or to maximize the Rényi of the compressive measurements.

The key to the approach considered in this paper is the realization that the projection-design problem for CS systems (subject to a power constraint) exhibits parallels with the precoder design problem for multiple-input–multiple-output (MIMO) communications systems: in the communications problem a source is being matched to a channel whereas in CS a channel, or equivalently the noise covariance, is being matched to the source. This link has also been recognized recently by Schnitter [43], who has provided projections designs for sources modelled by multivariate Gaussian distributions, as well as by Carson *et al.* [10], who have also provided designs for general multivariate source distributions. With the precoder design problem exhibiting a long tradition in the information theory and communications field, this link also provides the means to translate, with appropriate modifications, much of the design know-how and experience from the communications domain to CS.

The traditional problem of precoder design for MIMO Gaussian channels has been drawing on various performance metrics relevant for data communications. Common precoder design approaches aim to maximize the system signal-to-noise ratio (SNR) and the system signal-to-interference-plus-noise ratio (SINR) [33, 41] or minimize the system error probability [5], [18]. Another emerging precoder design approach imbued with operational significance is based on the maximization of the mutual information between the input and the output of the system [28, 34, 27, 36, 49]. This novel design principle has been shown to yield considerable rate gains in a variety of communications scenarios, due to the fact that, in addition to adapting to the channel characteristics, the designs also adapt to important features necessary to achieve high-rate reliable communications (the designs conform to the exact characteristics rather than only to the second-order statistics of the signaling scheme, as in traditional approaches (see [33], [41])). The basis of the emergence of the mutual information based designs have been fundamental connections between information theory and estimation theory, which have unveiled the interplay between mutual information and the minimum mean-squared error (MMSE) in scalar Gaussian channels [19] or mutual information and the MMSE matrix in vector Gaussian channels [32]. These results offer a means to bypass the absence of closed-form mutual information expressions for MIMO Gaussian channels driven by arbitrary (non-Gaussian) signaling schemes.

The operational significance of mutual information, which acts as the rationale for its use as the basis of a plethora of designs, is well known not only in data communications – it represents the highest reliable information transmission rate in a single-user channel driven by a specific signalling scheme – but also in other domains. For example, in classification problems mutual information relates (through bounds) to the Bayesian error probability of the classifier [21]; and, in regression problems

mutual information relates (through bounds) to the reconstruction error [38].

We consider design of the measurement kernel based upon maximizing the mutual information between the underlying signal \mathbf{x} and the compressive measurement \mathbf{y} . We also consider design based upon maximizing the Rényi entropy of \mathbf{y} , where the latter represents a generalization with operational relevance [16]. The projection design will be implemented in practice using gradient descent, and we demonstrate that for a GMM signal model the gradient of Rényi entropy with respect to the design matrix may be expressed analytically, for a special parameter setting. Further, we recover the gradient of Shannon entropy as a special case of the Rényi result.

The article considers both theoretical results, which disclose key operations effected by the projection designs, as well as experimental results that demonstrate the merit of the approach as applied to a practical CS imaging problem. One key operation relates to the notion of *mode alignment* in mutual information based designs: the modes of the source, which depending on the source statistical model are given by the eigenvectors of the source covariance matrix or the eigenvectors of the source MMSE matrix, have to align with the modes (eigenvectors) of the noise covariance matrix as a means to improve performance. This role can also be conceptually appreciated by viewing the measurement kernel as a sieve that aligns relevant statistical features of the source to the statistical features of the noise, in order to disclose relevant information for reconstruction. The relevance of mode alignment, which is typically absent in communications problems¹, has also been recently unveiled in radar applications [46].² Overbridging the theoretical and practical results is also the formal justification of a low-complexity high-performance online projections strategy, the partial direction sensing method (PDS) [14], which brings together the main operational features of the optimal measurement kernel designs, including mode alignment.

The detailed contributions of the article include:

- Recognition that recent advances in communications, which relate to the design of precoders for MIMO communications channels, carry over to CS, leading to a communications-inspired kernel design framework for CS applications.
- Proposal of mutual information based offline – where a set of projections is optimized simultaneously – and online – where the individual projections are optimized sequentially – kernel designs. The article unveils key operations carried out by the optimal kernel designs for multivariate Gaussian sources and general multivariate sources, including the operations of source and noise modes exposure, mode weighting and mode alignment. Particular emphasis is given to the role of mode alignment as a means to improve further the reconstruction performance in compressive sensing applications.
- Proposal of Rényi entropy based kernel designs. The article also underlines some relations between the mutual information (or Shannon entropy) and the Rényi entropy based kernel constructions.
- Formal rationale for the PDS strategy [14], which is based on the operational insight unveiled by the theoretical characterizations of the optimal kernel

¹This operation is absent in precoder designs for MIMO Gaussian channels driven by Gaussian inputs, due to the fact that the signal covariance is often taken to be white, but is present in precoder designs for MIMO Gaussian channels driven by non-Gaussian inputs. The role of a certain permutation operation in the precoder design is hinted at by Lamarca in [27].

²Note that Schniter [43] does not recognize the role of mode alignment due to the statistical assumptions about the source and noise covariances: this operation is not present when the source covariance is the identity matrix or when the noise covariance is also the identity matrix.

designs.

- Partial generalizations of the I-MMSE identity from the mutual information (or Shannon entropy) to the Rényi entropy domain.
- A range of experimental results that illustrate the benefit of the novel measurement kernel designs in relation to the conventional random ones.

The remainder of the article is organized as follows. In Section 2 we briefly summarize the notation used throughout. Section 3 reviews the modeling and design approach, introducing key system assumptions. Section 4 introduces the optimal kernel design based on the Shannon-based mutual information metric – this builds upon work on the communications field on precoder design for MIMO channels driven by Gaussian inputs and arbitrary inputs. Section 5 introduces the optimal kernel design based on the Rényi entropy metric, taking advantage of the closed-form expressions available for a GMM source. Section 6 provides the body of evidence that demonstrates the performance improvement possible through the application of the projections designs put forth in previous sections. We consider examples based on offline kernel design, based upon the prior signal model, as well as online kernel design based upon sequential update of the posterior, all within the context of a GMM signal representation, which yields analytic CS inversion. Section 7 draws the main conclusions. The Appendices contain proofs and supporting mathematical derivations.

2. Notation and Definitions. In the following text scalar quantities are denoted by italics, vectors are denoted by boldface lower case letters and matrices are denoted by boldface upper case letters. The projection of scalar x onto the non-negative orthant is denoted $(x)^+ \triangleq \max(0, x)$. The superscript $(\cdot)^*$ is used to denote an optimal solution and the superscripts $(\cdot)^T$, $(\cdot)^*$ and $(\cdot)^\dagger$ denote transpose, conjugate and conjugate transpose operators, respectively. The element in the i -th row and j -th column of the matrix \mathbf{X} is denoted by $[\mathbf{X}]_{i,j}$. The trace of a matrix is denoted $\text{tr}(\cdot)$. The diagonal matrix with diagonal elements given by either vector \mathbf{x} or the diagonal elements of matrix \mathbf{X} is denoted by $\text{Diag}(\mathbf{x})$ or $\text{Diag}(\mathbf{X})$, respectively.

We refer frequently to the following special matrices and sets: the $n \times n$ identity matrix is denoted \mathbf{I}_n , the $n \times n$ flipped identity matrix with ones on the anti-diagonal is denoted \mathbf{J}_n , the $m \times n$ matrix of all zeros is denoted $\mathbf{0}_{m \times n}$; the sub-scripts may be dropped where no confusion may arise. The set of all $n \times n$ unitary matrices is denoted \mathbb{S}^n and the set of $m \times n$ complex matrices is denoted $\mathbb{C}^{m \times n}$.

The notation $\mathbf{x} \sim \mathcal{CN}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a random variable \mathbf{x} which is circularly symmetric complex Gaussian distributed with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

3. Modelling and Design Approach. In CS, we aim to reconstruct the signal of interest $\mathbf{x} \in \mathbb{C}^m$ based on a small number of noisy projections:

$$\mathbf{y} = \mathbf{M} \mathbf{x} + \mathbf{w}, \quad \mathbf{y} \in \mathbb{C}^\ell \quad (3.1)$$

with $\ell \leq m$ and where $\mathbf{M} \in \mathbb{C}^{\ell \times m}$ is the kernel (or projection) matrix and \mathbf{w} represents zero-mean circularly symmetric complex Gaussian noise with positive definite covariance matrix $\boldsymbol{\Sigma}_{\mathbf{w}}$, i.e., $\mathbf{w} \sim \mathcal{CN}(\mathbf{w}; \mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{w}})$. The action of the kernel can be understood in terms of two separate projections and a power allocation (or stretching) operation, which are associated with the matrices in its singular value decomposition (SVD) given by:

$$\mathbf{M} = \mathbf{U}_{\mathbf{M}} \boldsymbol{\Lambda}_{\mathbf{M}} \mathbf{V}_{\mathbf{M}}^\dagger \quad (3.2)$$

where $\mathbf{\Lambda}_{\mathbf{M}} = [\text{Diag}(\sqrt{\lambda_{M_1}}, \dots, \sqrt{\lambda_{M_\ell}}) \quad \mathbf{0}_{\ell \times (m-\ell)}] \in \mathbb{R}^{\ell \times m}$, $\mathbf{U}_{\mathbf{M}} \in \mathbb{S}^\ell$, $\mathbf{V}_{\mathbf{M}} \in \mathbb{S}^m$, and $\lambda_{M_1} \geq \lambda_{M_2} \geq \dots \geq \lambda_{M_\ell} \geq 0$ correspond to the (non-negative) eigenvalues of $\mathbf{M} \mathbf{M}^\dagger$.

Both the signal and the noise covariance matrices are positive (semi-)definite and can also be represented in terms of their eigenvalue decomposition (projections and power allocation). In particular, the signal covariance matrix is given by:

$$\mathbf{\Sigma}_{\mathbf{x}} = \mathbf{U}_{\mathbf{x}} \mathbf{\Lambda}_{\mathbf{x}} \mathbf{U}_{\mathbf{x}}^\dagger \quad (3.3)$$

where $\mathbf{U}_{\mathbf{x}} \in \mathbb{S}^m$, $\mathbf{\Lambda}_{\mathbf{x}} = \text{Diag}(\lambda_{x_1}, \dots, \lambda_{x_m})$ and $\lambda_{x_1} \geq \lambda_{x_2} \geq \dots \geq \lambda_{x_m} \geq 0$ are the (non-negative) eigenvalues of $\mathbf{\Sigma}_{\mathbf{x}}$. Similarly, the noise covariance matrix is given by:

$$\mathbf{\Sigma}_{\mathbf{w}} = \mathbf{U}_{\mathbf{w}} \mathbf{\Lambda}_{\mathbf{w}} \mathbf{U}_{\mathbf{w}}^\dagger \quad (3.4)$$

where $\mathbf{U}_{\mathbf{w}} \in \mathbb{S}^\ell$, $\mathbf{\Lambda}_{\mathbf{w}} = \text{Diag}(\lambda_{w_1}, \dots, \lambda_{w_\ell})$ and $0 \leq \lambda_{w_1} \leq \lambda_{w_2} \leq \dots \leq \lambda_{w_\ell}$ correspond to the (non-negative) eigenvalues of $\mathbf{\Sigma}_{\mathbf{w}}$.

Our design approach, which relies not only on a statistical model for the noise but also on the signal, draws on specific quantitative metrics in order to conceive and compare various possible kernel designs. A natural metric, which relates to the best achievable reconstruction error, is the (non-linear) MMSE given by:

$$\text{MMSE} = \mathbb{E} \left\{ \text{tr} \left[(\mathbf{x} - \mathbb{E}\{\mathbf{x}|\mathbf{y}\}) (\mathbf{x} - \mathbb{E}\{\mathbf{x}|\mathbf{y}\})^\dagger \right] \right\} \quad (3.5)$$

that involves the use of conditional mean estimation to recover the signal of interest from the noisy projections, i.e., $\hat{\mathbf{x}}(\mathbf{y}) = \mathbb{E}\{\mathbf{x}|\mathbf{y}\}$ [25]. We, however, capitalize on information-theoretic metrics, most notably the mutual information and Rényi entropy based on the fact that mutual information and Rényi entropy - in view of recent developments in information theory and communications - appear to be more amenable to mathematical analysis than the non-linear MMSE. In addition, it is also possible to bound the MMSE via the mutual information as follows [38]:

$$\text{MMSE} \geq \frac{1}{2\pi e} \exp 2 [\mathcal{H}_x(\mathbf{x}) - \mathcal{I}(\mathbf{x}; \mathbf{y})]. \quad (3.6)$$

where $\mathcal{H}_x(\mathbf{x})$ denotes the differential entropy of \mathbf{x} and $\mathcal{I}(\mathbf{x}; \mathbf{y})$ denotes the mutual information between \mathbf{x} and \mathbf{y} .

The crux of our design approach, which we also partially extend from the mutual information to the Rényi entropy metric, is a fundamental result that links the gradient with respect to some parameters of the mutual information between the input and the output of a linear vector Gaussian channel model and the MMSE matrix associated with the model: known as the I-MMSE relationship. This result, which was originally put forth for the linear scalar Gaussian model by Guo, Shamai and Verdú [19] and later for linear vector Gaussian channels by Palomar and Verdú in [32], can be directly applied to the model in (3.1) so that:

$$\nabla_{\mathbf{M}} \mathcal{I}(\mathbf{x}; \mathbf{y}) = \mathbf{\Sigma}_{\mathbf{w}}^{-1} \mathbf{M} \mathbf{E} \quad (3.7)$$

where the MMSE matrix is³:

$$\mathbf{E} = \mathbb{E} \left\{ (\mathbf{x} - \mathbb{E}\{\mathbf{x}|\mathbf{y}\}) (\mathbf{x} - \mathbb{E}\{\mathbf{x}|\mathbf{y}\})^\dagger \right\} \quad (3.8)$$

$$= \mathbf{U}_{\mathbf{E}} \mathbf{\Lambda}_{\mathbf{E}} \mathbf{U}_{\mathbf{E}}^\dagger \quad (3.9)$$

³Note that the MMSE matrix \mathbf{E} is a function of the kernel \mathbf{M}

where $\mathbf{U}_{\mathbf{E}} \in \mathbb{S}^m$, $\mathbf{\Lambda}_{\mathbf{E}} = \text{Diag}(\lambda_{E_1}, \dots, \lambda_{E_m})$ and $\lambda_{E_1} \geq \lambda_{E_2} \geq \dots \geq \lambda_{E_m} \geq 0$ are the (non-negative) eigenvalues of the MMSE matrix.

Our design approach also draws on a specific kernel design constraint. It is important to recall that in CS applications the kernel design is typically set to obey unit-norm row constraints or, instead, orthonormal constraints [30]. In contrast, in communications applications the kernel (or precoder) design obeys a power (trace) constraint, which states that on average the rows have unit-norm. A paper that does consider this power constraint for CS is the work by Schnitter [43], however, this is unusual in the CS field. We adopt this more general constraint, which, in addition to leading to solutions with higher mutual information or Rényi entropy, enables the formulation of the design framework which the unit-norm rows constraint does not. The exception to this is the special case of adaptive online design in Section 6.2, where each row of the kernel is designed sequentially such that the two constraints coincide.

4. Mutual Information based Kernel Design. In this section we consider the characterization of the kernel that maximises the mutual information of the model in (3.1), subject to a power constraint, for multivariate Gaussian sources and general multivariate sources. The optimal kernel design for multivariate Gaussian sources also provides a rationale for other kernel designs in subsequent sections, most notably, the PDS method (we extend the work of [14]). The design problem can be posed abstractly as follows:

$$\begin{aligned} & \underset{\mathbf{M}}{\text{maximize}} && \mathcal{I}(\mathbf{x}; \mathbf{M} \mathbf{x} + \mathbf{w}) \\ & \text{subject to} && \frac{1}{\ell} \text{tr}(\mathbf{M}\mathbf{M}^\dagger) \leq 1 \end{aligned} \tag{4.1}$$

It is important to remark that this optimization problem is non-convex in general. The use of the fundamental result in (3.7), in addition to enabling the full or partial characterization of the solution, also leads to efficient computational procedures. We restate next the characterizations of the optimal kernel designs for Gaussian sources (Theorem 4.1) and general sources (Theorem 4.2), which also appear in slightly different forms in [36] [35] [27], in a manner that emphasizes the operational significance for CS applications.

4.1. Multivariate Gaussian Input Source. The characterization of the optimal kernel design for a multivariate complex-valued Gaussian source leverages the well-known closed-form mutual information expression given by:

$$\mathcal{I}(\mathbf{x}; \mathbf{y}) = \log \det (\mathbf{I}_m + \mathbf{M}^\dagger \mathbf{\Sigma}_{\mathbf{w}}^{-1} \mathbf{M} \mathbf{\Sigma}_{\mathbf{x}}). \tag{4.2}$$

This simple closed-form expression allows the use of simple matrix identities, rather than the gradient result in (3.7), to obtain the solution to (4.1). The case when $\mathbf{\Sigma}_{\mathbf{x}} = \mathbf{I}$ is well-known from communications theory and was recently applied in the design of measurement kernels by Schnitter [43]. However, the case for general source covariance matrices has not been studied in the communications domain. We unveil that this leads to the novel operation of *mode alignment*⁴.

THEOREM 4.1. *The kernel matrix that solves the optimization problem in (4.1) for a multivariate complex-valued Gaussian source with covariance matrix $\mathbf{\Sigma}_{\mathbf{x}}$ is given*

⁴This result was also recently shown in radar [46].

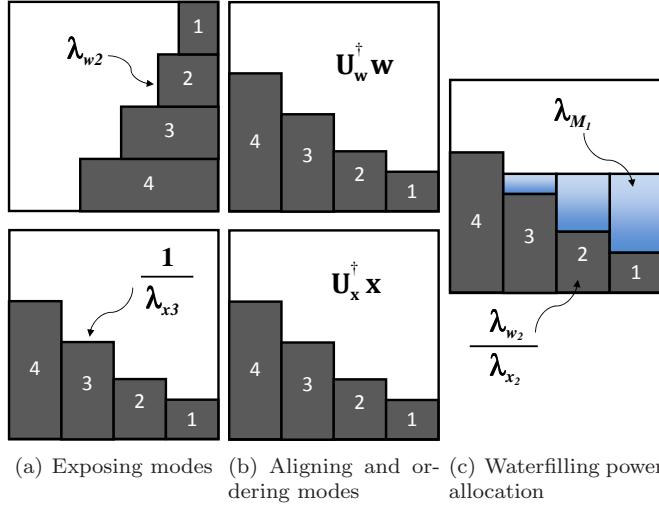


FIG. 4.1. Diagrammatic view of the actions of the optimal kernel design.

by⁵:

$$\mathbf{M}^* = \mathbf{U}_w \mathbf{\Lambda}_M^* \mathbf{U}_x^\dagger \quad (4.3)$$

where $\mathbf{\Lambda}_M^* = [\text{Diag}(\sqrt{\lambda_{M_1}^*}, \dots, \sqrt{\lambda_{M_\ell}^*}) \quad \mathbf{0}_{\ell \times (m-\ell)}]$, $\lambda_{M_i}^* = \left(\frac{1}{\eta} - \frac{\lambda_{w_i}}{\lambda_{x_i}}\right)^+$ with the noise covariance eigenvalues λ_{w_i} arranged in ascending order and the source covariance eigenvalues λ_{x_i} arranged in descending order and η ensures the average unit-norm row constraint, i.e., $\frac{1}{\ell} \text{tr}(\mathbf{M}\mathbf{M}^\dagger) = 1$.

Proof. See Appendix C. \square

Theorem 4.1 uncovers the operations of the optimal kernel design. In particular, it is possible to recognize a novel mode alignment operation which involves two aspects: i) exposing the modes of the noise and source covariance; and ii) ordering (or aligning) the modes.

First, the left-singular vectors of the kernel are chosen to align with the eigenvectors of the noise covariance matrix and the right-singular vectors of the kernel are chosen to align with the eigenvectors of the signal covariance matrix (Fig. 4.1(a)). This is referred to as exposing the modes.

The ordering (or alignment) of the exposed modes is very particular, the largest source eigenvalue is matched to the smallest noise eigenvalue, the second largest source eigenvalue is matched to the second smallest noise eigenvalue, and so on (Fig. 4.1(b)).

Finally, the kernel “weights” the matched modes according to a “waterfilling” interpretation [12] (Fig. 4.1(c)). Intuitively, this emphasizes the less noisy “channels” and reduces the influence of the noisier ones as a means to maximize further mutual information.

As an example, Fig. 4.2 depicts the mutual information associated with two possible alignments for the signal and noise eigenvalues in a scenario where both

⁵Note that the superscript \star denotes an optimal solution.




FIG. 4.2. Mutual information as a function of SNR for two different alignments both with optimal power allocation.

covariance matrices are diagonal, $\mathbf{\Lambda}_{\mathbf{x}} = \text{Diag}(1, 0.25)$ and $\mathbf{\Lambda}_{\mathbf{w}} = \text{Diag}(1, 0.25)$. It is evident that the ordering of the modes has a significant impact on the mutual information at low and medium SNR – the highest mutual information corresponds to the kernel design that aligns the strongest source eigenvalue with the weakest noise eigenvalue, $\mathbf{U}_{\mathbf{M}} = \mathbf{U}_{\mathbf{w}} = \mathbf{J}_2$ and $\mathbf{V}_{\mathbf{M}} = \mathbf{I}_2$.

4.2. General Multivariate Input Source. While the application of communications theory results for Gaussian distributions are known to varying degrees outside the field of communications theory, the results for general sources have not been fully leveraged outside of communications. The characterization of the optimal kernel design for a general multivariate complex-valued source, in view of the absence of closed-form mutual information expressions, now leverages the fundamental result in (3.7).

THEOREM 4.2. *The kernel matrix that solves the optimization problem in (4.1) for a general multivariate complex-valued source with covariance matrix $\mathbf{\Sigma}_{\mathbf{x}}$ is given by:*

$$\mathbf{M}^* = \mathbf{U}_{\mathbf{w}} \mathbf{\Lambda}_{\mathbf{M}}^* \mathbf{V}_{\mathbf{M}}^{*\dagger} \quad (4.4)$$

where $\mathbf{V}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}^*$ ⁶, the matrix $\mathbf{\Pi}^*$ is the optimal permutation matrix, $\mathbf{\Lambda}_{\mathbf{M}}^* = [\text{Diag}(\sqrt{\lambda_{M_1}^*}, \dots, \sqrt{\lambda_{M_\ell}^*}) \quad \mathbf{0}]$, and $\lambda_{M_i}^*$ are given by the generalized mercury water-filling solution, i.e.,

$$\lambda_{M_i}^* = \begin{cases} 0, & \eta \lambda_{w_i} > \text{mmse}_i(\mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}^*, \mathbf{\Lambda}_{\mathbf{Q}}^* |_{\lambda_{M_i}^*=0}) \\ \text{mmse}_i^{-1}(\eta \lambda_{w_i}), & \text{otherwise} \end{cases} \quad (4.5)$$

where η ensures the average unit-norm row constraint, i.e., $\frac{1}{\ell} \text{tr}(\mathbf{M}^* \mathbf{M}^{*\dagger}) = 1$, $\mathbf{\Lambda}_{\mathbf{Q}}^* = \mathbf{\Lambda}_{\mathbf{M}}^* \mathbf{\Lambda}_{\mathbf{M}}^{*\dagger}$, $\mathbf{\Lambda}_{\mathbf{Q}}^* |_{\lambda_{M_i}^*=0} = \text{Diag}(\lambda_{M_1}^*, \dots, \lambda_{M_{i-1}}^*, 0, \lambda_{M_{i+1}}^*, \dots, \lambda_{M_\ell}^*)$ and $\text{mmse}_i(\mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}^*, \mathbf{\Lambda}_{\mathbf{Q}}^*)$

⁶Note that the MMSE matrix \mathbf{E} is a function of the kernel \mathbf{M} .

denotes the i -th diagonal entry of the MMSE matrix associated with the estimate of $\mathbf{x}' = \mathbf{\Pi}^{\star\dagger} \mathbf{U}_{\mathbf{E}}^{\star\dagger} \mathbf{x}$ from

$$\mathbf{y}' = \mathbf{\Lambda}_{\mathbf{w}}^{-1/2} \mathbf{\Lambda}_{\mathbf{M}}^{\star} \mathbf{x}' + \mathbf{n}, \quad (4.6)$$

where \mathbf{n} is zero-mean circularly symmetric complex Gaussian noise with identity covariance, $\mathbf{n} \sim \mathcal{CN}(\mathbf{n}; \mathbf{0}, \mathbf{I})$. Note that mmse_i^{-1} is the inverse of mmse_i with respect to λ_{M_i} for fixed $\lambda_{M_j}, \forall j \neq i$.

Proof. See Appendix D. \square

REMARK 1. In the high noise/low signal power regime, a first-order expansion of the mutual information is given [32]:

$$\mathcal{I}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \text{tr}(\mathbf{\Sigma}_{\mathbf{w}}^{-1} \mathbf{M} \mathbf{\Sigma}_{\mathbf{x}} \mathbf{M}^{\dagger}) + o(\|\mathbf{\Sigma}_{\mathbf{x}}\|) \quad (4.7)$$

which implies the result observed by Shannon that at low signal-to-noise ratios proper complex discrete inputs offer a negligible loss in performance terms with regards to the capacity achieved by Gaussian inputs; hence the results for the Gaussian in Theorem 4.1 also apply in general for proper complex sources in the high noise/low power regime.

Theorem 4.2 suggests that the *mode alignment* is no longer between the eigenvectors of the source covariance and the eigenvectors of the noise covariance, but between the eigenvectors of the MMSE matrix and the eigenvectors of the noise covariance. The diagonalization of the MMSE matrix was first noted for communications by Lamarca [27] for identity source covariances, and the same holds true for CS for general source covariances. The singular values of the kernel are described by the mercury waterfilling algorithm [28] [36] which differs from waterfilling by adjusting for the non-Gaussian nature of the inputs, however, the procedure is remarkably similar.

It is important to emphasize that Theorem 4.1 characterizes fully the optimal kernel design but - in view of the non-convexity of the problem - Theorem 4.2 characterizes partially, via a fixed point equation, the optimal kernel since $\mathbf{U}_{\mathbf{E}}^{\star}$ is still a function of \mathbf{M} . The characterization is useful because it leads to i) stopping criteria for gradient descent algorithms via (3.7); and ii) alternative optimization algorithms. Note that if we implement gradient descent with (3.7) we may get trapped in local maxima since it is known that the mutual information is not always a concave function of \mathbf{M} [34]. However, mutual information is known to be concave in the squared singular values of \mathbf{M} , for $\mathbf{U}_{\mathbf{M}} = \mathbf{U}_{\mathbf{w}}$ and fixed $\mathbf{V}_{\mathbf{M}}$. An alternative gradient descent algorithm that leads to the global maximum by avoiding local maxima switches between optimizing the singular values and the right-singular vectors of the kernel [49].

5. Design with Rényi Entropy. We consider the characterization of the kernel that maximizes the output Rényi entropy of the model in (3.1), subject to a power constraint, for multivariate Gaussian sources and multivariate Gaussian mixture sources. The design problem can then be posed as follows:

$$\begin{aligned} & \underset{\mathbf{M}}{\text{maximize}} && h_{\alpha}(\mathbf{M} \mathbf{x} + \mathbf{w}) \\ & \text{subject to} && \frac{1}{\ell} \text{tr}(\mathbf{M} \mathbf{M}^{\dagger}) \leq 1 \end{aligned} \quad (5.1)$$

where

$$h_\alpha(\mathbf{y}) = \frac{1}{1-\alpha} \log \int p^\alpha(\mathbf{y}) d\mathbf{y}. \quad (5.2)$$

Note that Rényi entropy represents a generalization of Shannon entropy given by:

$$h_s(\mathbf{y}) = - \int p(\mathbf{y}) \log p(\mathbf{y}) d\mathbf{y} \quad (5.3)$$

which is the special case when $\alpha = 1$.

5.1. Multivariate Gaussian Input Source. For multivariate Gaussian sources, both Shannon entropy and Rényi entropy can be expressed analytically for all values of $\alpha > 0$. In particular, the two are shown to be related in the following theorem⁷:

THEOREM 5.1.

For a multivariate Gaussian input source where $\mathbf{x} \sim \mathcal{CN}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}_{\mathbf{x}})$, the Rényi entropy of order $\alpha > 0$ and the Shannon entropy associated with the output of the model in (3.1) are related as:

$$h_\alpha(\mathbf{y}) = h_s(\mathbf{y}) - \ell \left(1 - \frac{\log \alpha}{\alpha - 1} \right), \quad (5.4)$$

where $h_s(\mathbf{y}) = \log [(2\pi e)^\ell \det (\boldsymbol{\Sigma}_{\mathbf{w}} + \mathbf{M} \boldsymbol{\Sigma}_{\mathbf{x}} \mathbf{M}^\dagger)]$.

Proof. See Appendix E. \square

Theorem 5.1 leads immediately to a generalization of the I-MMSE identity in (3.7) for Gaussian sources:

THEOREM 5.2. *For Gaussian sources, the (complex) gradient with respect to the kernel of the output Rényi entropy of order $\alpha > 0$ associated with the model in (3.1) obeys the relationship:*

$$\nabla_{\mathbf{M}} h_\alpha(\mathbf{y}) = \boldsymbol{\Sigma}_{\mathbf{w}}^{-1} \mathbf{M} \mathbf{E}. \quad (5.5)$$

Theorem 5.2 unveils that the relationship between mutual information and the MMSE matrix in (3.7) also holds for all values of $\alpha > 0$ for the output Rényi entropy associated with the model in (3.1) for Gaussian sources. Theorem 5.2 also implies that the kernel design that maximizes the Rényi entropy subject to a power constraint also obeys the characterization in Theorem 4.1.

5.2. Multivariate Gaussian Mixture Model Input Source. For Gaussian Mixture Models (GMM) the signal $\mathbf{x} \in \mathbb{C}^m$ is represented by:

$$p(\mathbf{x}) = \sum_{i=1}^N p(i) \mathcal{CN}(\mathbf{x}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i), \quad (5.6)$$

where $p(i)$ is the probability of occurrence of mixture component i , $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ correspond to the mean and covariance matrix of the i -th circularly symmetric complex

⁷For quadratic Rényi entropy this result was also in Appendix A of [42].

Gaussian distribution. Neither the Shannon entropy, mutual information nor the MMSE matrix are known to have closed-form expressions for GMMs. Rényi entropy and its gradient, however, admit closed-form expressions in some instances, which lend themselves more easily to optimization via gradient descent algorithms. For example, the quadratic Rényi entropy of the noisy projection \mathbf{y} in (3.1) is given by:

$$h_2(\mathbf{y}) = -\log \sum_{i=1}^N \sum_{j=1}^N p(i) p(j) \mathcal{CN}(\mathbf{0}; \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j}) \quad (5.7)$$

where:

$$\boldsymbol{\mu}_{i,j} = \mathbf{M}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \quad (5.8)$$

$$\boldsymbol{\Sigma}_{i,j} = \mathbf{M}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) \mathbf{M}^\dagger + 2\boldsymbol{\Sigma}_{\mathbf{w}} \quad (5.9)$$

The complex gradient with respect to \mathbf{M} of the quadratic Rényi entropy of the noisy projection \mathbf{y} in (3.1) for the GMM is given by:

$$\begin{aligned} \nabla_{\mathbf{M}} h_2(\mathbf{y}) = & \\ & - \frac{\sum_{i,j=1}^N p(i) p(j) \mathcal{CN}(\mathbf{0}; \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j}) \nabla_{\mathbf{M}} \log \mathcal{CN}(\mathbf{0}; \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j})}{\sum_{i=1}^N \sum_{j=1}^N p(i) p(j) \mathcal{CN}(\mathbf{0}; \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j})} \end{aligned} \quad (5.10)$$

where:

$$\begin{aligned} \nabla_{\mathbf{M}} \log \mathcal{CN}_{i,j} = & -\boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) \\ & + \boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\dagger \\ & \times \{\mathbf{M}^\dagger \boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) - \mathbf{I}\}. \end{aligned} \quad (5.11)$$

where \times denotes a matrix multiplication. The proof is given in Appendix F.

It is interesting to note that the now celebrated I-MMSE relationship in the information theory literature also applies for Rényi entropy of order $\alpha > 0$ associated with Gaussian source models. However, this relationship does not seem to carry over for the Rényi entropy of more general source models. In fact, it can only be shown that for a general source, which obeys some additional smoothness conditions, the gradient can be expressed as follows (The proof is a modification of the result in [32]):

$$\nabla_{\mathbf{M}} h_\alpha(\mathbf{y}) = \alpha \boldsymbol{\Sigma}_{\mathbf{w}}^{-1} \int \hat{p}(\mathbf{y}) (\mathbf{y} - \mathbf{M}\mathbf{x}_y) \mathbf{x}_y^\dagger d\mathbf{y} \quad (5.12)$$

where the probability distribution $\hat{p}(\mathbf{y}) = \frac{p^\alpha(\mathbf{y})}{\int p^\alpha(\mathbf{y}) d\mathbf{y}}$ and \mathbf{x}_y is the conditional mean estimator.

It is not difficult to appreciate that the right-hand side of (5.12) is in general different from the right-hand side of (5.5) (or the right-hand side of the I-MMSE relationship in (3.7)) by studying the Taylor expansion of $\nabla_{\mathbf{M}} h_2(\mathbf{y})$. For the high-noise power scenario, the first term in the expressions coincide but higher order terms do not [39].

6. Application to Compressive Sensing.

6.1. Problem setup. We consider CS in the context of imaging. While the theory is applicable to complex data, the following examples focus on real images. Specifically, consider measurement of the image $\mathbf{X} \in \mathbb{R}^{N_x \times N_y}$, for large N_x and N_y . As indicated in Figure 6.1, the image is partitioned into $n_x \times n_y$ contiguous “patches,” with the pixels in the j th patch denoted by vector $\mathbf{x}_j \in \mathbb{R}^\ell$, with $\ell = n_x n_y$. In the examples considered here $n_x = n_y = 8$ (consistent, for example, with the patch sizes used in the JPEG standard).

It is desirable to partition the images into such patches because one may readily learn a signal model for the $\{\mathbf{x}_j\}$, while it is difficult to learn an accurate signal model directly on the entire image \mathbf{X} . Specifically, following [11, 14], we assume that each \mathbf{x}_j is drawn from a GMM of the form (5.6), here for *real* normal distributions.

To learn the prior signal model $p(\mathbf{x})$ for the patches, we first consider a large ensemble of natural images, from which patches $\mathbf{x}_j \in \mathbb{R}^\ell$ are selected at random. Using these training data, a (real) GMM of the form in (5.6) is constituted as a signal model. To learn this GMM, we have employed nonparametric Bayesian methods as in [11], as well as expectation-maximization (EM) methods [14], and both methods yield very similar results. The following results are based on a $N = 20$ component GMM, trained on 100,000 patches, extracted at random from 500 natural images in the Berkeley Segmentation Dataset (<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/resources.html>). These training images are distinct from those considered in the testing phase, for CS inversion.

While patches are selected at random from training images to constitute the prior $p(\mathbf{x})$, when performing CS the goal is to recover the entire underlying image \mathbf{X} . Therefore, for CS inversion we wish to recover each of the $\{\mathbf{x}_j\}$ in Figure 6.1. In general, a separate projection matrix \mathbf{M}_j is applied to patch j from image \mathbf{X} . For the case of offline design of the projection matrix, \mathbf{M}_j is the same for all patches j (since it is non-adaptive). For online design a distinct \mathbf{M}_j is adaptively designed for each testing patch j . The measured data associated with patch j is expressed as

$$\mathbf{y}_j = \mathbf{M}_j \mathbf{x}_j + \mathbf{w}_j, \quad j = 1, \dots, J \quad (6.1)$$

In the examples that follow, the images under test are 256×256 , and therefore this procedure was employed on $J = 1024$ non-overlapping patches of size 8×8 . Each of the \mathbf{x}_j are recovered independently from the respective measured \mathbf{y}_j , thereby allowing for massive parallelization.

For simplicity, we henceforth drop the subscript j , and the discussion that follows applies to each of the J patches in Figure 6.1. We assume the noise $\mathbf{w} \sim \mathcal{N}(0, \Sigma_{\mathbf{w}})$, with known covariance matrix $\Sigma_{\mathbf{w}}$. In the following examples we consider low-noise, i.i.d. measurements, and therefore $\Sigma_{\mathbf{w}} = 10^{-6} \mathbf{I}_\ell$. The likelihood function for the underlying signal \mathbf{x} is $\mathcal{N}(\mathbf{y}; \mathbf{M}\mathbf{x}, \Sigma_{\mathbf{w}})$, and the prior $p(\mathbf{x})$ is the aforementioned GMM, $p(\mathbf{x}) = \sum_{i=1}^N w_i \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_i, \Sigma_i)$. Under this likelihood function for \mathbf{x} , and with the GMM prior, the posterior $p(\mathbf{x}|\mathbf{y})$ is also a GMM:

$$p(\mathbf{x}|\mathbf{y}) = \sum_{i=1}^N \tilde{w}_i \mathcal{N}(\mathbf{x}; \tilde{\boldsymbol{\mu}}_i, \tilde{\Sigma}_i) \quad (6.2)$$

with

$$\tilde{\Sigma}_i^{-1} = \mathbf{M}^T \Sigma_{\mathbf{w}}^{-1} \mathbf{M} + \Sigma_i^{-1}, \quad \tilde{\boldsymbol{\mu}}_i = \tilde{\Sigma}_i (\mathbf{M}^T \Sigma_{\mathbf{w}}^{-1} \mathbf{y} + \Sigma_i^{-1} \boldsymbol{\mu}_i) \quad (6.3)$$




FIG. 6.1. Spatial grid used for CS measurement image $N_x \times N_y$ image, decomposing the image into a contiguous grid of “patches,” each patch composed of $n_x \times n_y$ pixels, $n_x \ll N_x$ and $n_y \ll N_y$. Letting $\mathbf{x}_i \in \mathbb{R}^{n_x n_y}$ represent the pixels associated with the i th patch, separate projection matrices \mathbf{M}_i are designed for each \mathbf{x}_i .

$$\tilde{w}_i = w_i \mathcal{N}(\mathbf{y}; \mathbf{M}\boldsymbol{\mu}_i, \mathbf{M}\boldsymbol{\Sigma}_i\mathbf{M}^T + \boldsymbol{\Sigma}_{\mathbf{w}}) / p(\mathbf{y}) \quad (6.4)$$

$$p(\mathbf{y}) = \sum_{i=1}^N w_i \mathcal{N}(\mathbf{y}; \mathbf{M}\boldsymbol{\mu}_i, \mathbf{M}\boldsymbol{\Sigma}_i\mathbf{M}^T + \boldsymbol{\Sigma}_{\mathbf{w}}) \quad (6.5)$$

When presenting results, the estimated signal $\hat{\mathbf{x}}$ is the mean based on $p(\mathbf{x}|\mathbf{y})$, *i.e.*, $\hat{\mathbf{x}} = \sum_{i=1}^N \tilde{w}_i \tilde{\boldsymbol{\mu}}_i$.

6.2. Offline and online design. We consider online and offline design of the projection matrix \mathbf{M} , based upon gradient descent: $\mathbf{M} \leftarrow \mathbf{M} + \gamma \nabla_{\mathbf{M}} \mathcal{I}(\mathbf{x}; \mathbf{y})$, with re-normalization to satisfy the power constraint; here we perform a gradient of the mutual information, and the same type of gradient descent is performed in the context of Rényi entropy, for which we therefore employ the results of Section IV. When employing the gradient of Rényi entropy, we employ (3.7). The design of \mathbf{M} based upon a gradient of mutual information is denoted PV, for Palomar and Verú.

For offline PV design, $p(\mathbf{x})$ corresponds to the learned prior GMM, and the entire \mathbf{M} is inferred at once. For online PV design, after measuring the first k components of \mathbf{y} , denoted $\mathbf{y}_{1:k}$, we update the posterior $p(\mathbf{x}|\mathbf{y}_{1:k})$ via (6.2), and row $k+1$ of \mathbf{M} is constituted based upon this posterior signal model; after each measurement, the posterior is updated, followed by design of the next row of \mathbf{M} , used to define the next measurement. In these computations, the MMSE matrix in (3.8) is computed via Monte Carlo integration, based on draws from $p(\mathbf{x})$ (in the offline case) or $p(\mathbf{x}|\mathbf{y}_{1:k})$ (in the online case). Online design of the patch-dependent projection matrix \mathbf{M} may be performed in parallel.



For Rényi-based design we consider the case $\alpha = 2$; this is convenient, as within the context of the GMM representation employed here the gradient with respect to \mathbf{M} is analytic, via (5.10).

The online PV design is relatively expensive, as one must repeatedly perform Monte Carlo integration to update the MMSE matrix \mathbf{E} , and one must also perform gradient descent. For online Rényi-based design we employ (5.10); while this analytic expression precludes the need to numerically compute \mathbf{E} , the large number of sums makes online Rényi and online PV design comparably expensive.

The relative expense of Rényi and PV online design motivates a simplified online design. In [14] the authors proposed the PDS method, in which a GMM was used for $p(\mathbf{x})$. In [14], the components of the first $k < \ell$ rows of \mathbf{M} are drawn i.i.d. from a zero-mean normal distribution. Using this k -row sensing matrix, an initial measurement $\mathbf{y}_{1:k} \in \mathbb{R}^k$ is performed. Based upon $\mathbf{y}_{1:k}$, the most probable mixture component from the prior $p(\mathbf{x})$ is selected. At this point a single-Gaussian signal model is constituted. The remaining $\ell - k$ rows of \mathbf{M} are then defined by the principal $\ell - k$ eigenvectors of the covariance matrix from this Gaussian. While [14] did not have access to our Theorem 1, the design so constituted is consistent with it. Specifically, Theorem 1 applies to the case of a single-Gaussian signal model. Under the aforementioned assumptions for $\Sigma_{\mathbf{w}}$ (diagonal covariance matrix, with small diagonal variance), Theorem 1 implies that the optimal projection matrix corresponds to the principal eigenvectors of the covariance matrix. However, the assumption of k initial random projections employed in [14], before selecting a single Gaussian component, seems undesirable. Further, in [14] the single Gaussian was selected from the prior $p(\mathbf{x})$ rather from the updated posterior $p(\mathbf{x}|\mathbf{y}_{1:k})$.

We extend the PDS technique to an online setting as follows. We first initialize $p(\mathbf{x})$ with the GMM prior signal model (learned using offline training data). We then sequentially constitute one row of \mathbf{M} at a time, from $k = 1, \dots, \ell$; after each row is so constituted, a single new projection measurement is performed with that new row. Again let $\mathbf{y}_{1:k}$ represent the vector of data constituted in this manner via the first k rows of \mathbf{M} . Based upon these data we update the signal model $p(\mathbf{x}|\mathbf{y}_{1:k})$. To design row $k+1$ of \mathbf{M} , let $i' = \arg\max_i \tilde{w}_i$, where the $\{\tilde{w}_i\}$ are the GMM mixture weights from $p(\mathbf{x}|\mathbf{y}_{1:k})$. Then the $(k+1)$ th row of \mathbf{M} is defined by the leading eigenvector of $\tilde{\Sigma}_{i'}$. The online PDS approximates the posterior GMM at each step with the dominant Gaussian from the posterior GMM $p(\mathbf{x}|\mathbf{y}_{1:k})$, and then via Theorem 1 the next row of \mathbf{M} is defined by the leading eigenvector of the associated covariance matrix. Since no Monte Carlo simulation and gradient descent are needed in the above process, online PDS method is very fast. The eigenvectors are orthonormal, and therefore the power constraint is satisfied automatically at every step. Note that the posterior $p(\mathbf{x}|\mathbf{y}_{1:k})$ continuously updates with increasing data, and therefore it is not particularly sensitive to the prior $p(\mathbf{x})$; the original PDS in [14] was based upon the prior $p(\mathbf{x})$ only, which may necessitate more care in selection of the training data. Since the posterior can be updated easily via (6.2), it appears highly preferable to use this approach rather than fixing the signal model.

6.3. Experimental Results. In Figures 6.2-6.4 results are shown for three widely examined test images: ‘barbara’, ‘house’ and ‘pepper,’ respectively. Two classes of results are considered based upon random projection design. The “random GMM” results employ the patch-based CS construction in Figure 6.1, and the learned GMM-based prior $p(\mathbf{x})$. The form of these results are the same as employed for the designed \mathbf{M}_j , except here each \mathbf{M}_j is constituted with matrix elements drawn

FIG. 6.2. *PSNR for the reconstructed 'barbara' image.*FIG. 6.3. *PSNR for the reconstructed 'house' image.*

i.i.d. from $\mathcal{N}(0,1)$, followed by normalization. We also considered CS design in which the projections are performed directly on the entire image \mathbf{X} , rather than at the patch level, as in Figure 6.1. If one performs CS inversion based on traditional CS algorithms, which employ ℓ_1 and related regularization [8], the quality of the inversion is markedly worse than that using the proposed approach, with learned signal models $p(\mathbf{x})$; we therefore do not show these results here, because they don't fit on the same scale of results presented. This is not surprising, as the patch-dependent learned signal model $p(\mathbf{x})$ is much richer, and tailored to the data than simple sparsity constraints, which motivate ℓ_1 regularization. To provide a fairer comparison, when performing inversion for the case in which the projections are performed directly on the entire image \mathbf{X} , we consider an underlying wavelet basis and perform inversion




FIG. 6.4. PSNR for the reconstructed ‘pepper’ image.

based on the sophisticated hidden Markov tree (HMT) wavelet model for images [20]. This signal model $p(\mathbf{x})$ could in principle also be used within the theory to design a projection matrix applicable to the entire image. However, the significant advantage of the GMM construction is that the posterior of the underlying signal may be constituted analytically, while for the HMT expensive computational methods are needed [20]. Therefore, we only show HMT inversion results when the projection matrix is constituted at random, thereby providing a comparison of inversion quality of the GMM (patch based) and the HMT (entire image), based upon random projections.

We consider offline design of the patch-based projection matrix \mathbf{M} based upon the Rényi measure of entropy, as well as based upon mutual information (via the PV theory). For online Rényi and PV design, we do *not* make a simplifying single-Gaussian assumption when designing each row of \mathbf{M} . By contrast the online PDS method uses the most probable Gaussian from the posterior to design the next projection at each step (this is therefore an approximation). The PDS method is very fast, while online PV is expensive, and therefore is shown principally for comparison (may not be done in practice, where online design must be fast).

First comparing the results based on random projections, the results based upon the (learned) patched-based GMM and based on the entire-image-based HMT are comparable in reconstruction quality. Sometimes the GMM results are slightly better, and other times the HMT results are better. However, there is no comparison with respect to computation speed. The HMT results are expensive, being based upon a Gibbs sampler [20]. By contrast the GMM results are very fast, with the inversion analytic. The additional big advantage of the GMM representation is that it allows convenient design of patch-dependent projection matrices, which we consider next.

Each of the designed projection methods yield significant improvement relative to random, and after approximately 6 projections per patch we note that the online results are significantly better than offline design. For the first approximately 5 measurements per patch, the offline and online results are comparable; we attribute this to an inadequate number of measurements to obtain an accurate signal model, and therefore little gain manifested by adaptivity. However, after approximately 6

measurements per patch it appears that the posterior signal model becomes accurate, yielding advantages of adaptivity. Concerning online design, inversion quality based on the simple and fast online PDS performs quite competitively relative to the on-line Rényi and PV design (which do not make a simplification to a single Gaussian), despite the fact that it assumes that the patch is drawn from a single Gaussian.

To understand the quality of the simple PDS-based design, consider Figure 6.5, wherein we plot the probabilities $\{\tilde{w}_i\}_{i=1,N}$, for the posterior $p(\mathbf{x}|\mathbf{y}_{1:k})$, as the number of measurements k increases from 1 to ℓ . Note that after approximately six measurements the model has inferred that the underlying signal \mathbf{x} was drawn from a single multivariate Gaussian. Note that the GMM is characteristic of an *ensemble* of draws, like those characteristic of the multiple patches in Figure 6.1. However, any single patch is drawn from a single one of the mixture components; it is however unknown *a priori* which component. Based upon experiments of this type, typically 6 projections are sufficient to infer which single mixture component a given patch corresponds to. At this point the results in Theorem 1 apply directly, which under the assumption for $\Sigma_{\mathbf{w}}$ dictates that the optimal measurement corresponds to projecting onto the dominant eigenvector of the covariance matrix of the single mixture component (single Gaussian); this is precisely what PDS does.




FIG. 6.5. Evolution of the mixture weight in the posterior GMM for a typical testing patch in ‘barbara’.

7. Conclusions. We observe that the design principle of maximizing mutual information or Rényi entropy leads to deterministic kernel matrices for which MMSE performance is superior to that of random kernel matrices. In particular, we are able to provide design principles for the optimal kernel matrix for a general multivariate source that maximizes the mutual information or Rényi entropy (for which Shannon entropy is a special case). We showed that the optimal kernel exposes the modes of the noise and the modes of the (optimal) MMSE matrix, then performs an alignment operation whose purpose it to optimally match the modes of the noise to the modes of the MMSE matrix (or, in the multivariate Gaussian source scenario, the modes of the source covariance). Finally, it carries out a generalized mercury-waterfilling power allocation operation.

The theoretical framework has been demonstrated with application to compressive sensing (CS) as applied to imagery. Using a GMM signal model, it was demonstrated that designed measurement kernels can yield markedly improved CS signal recovery relative to random design. The GMM representation has the advantage of yielding closed-form CS inversion, which is particularly attractive for fast signal inversion and for online kernel design. We have enhanced an online kernel design framework first proposed in [14], and have also provided a theoretical foundation for why it works so effectively in practice.

REFERENCES

- [1] A. ASHOK, P. BAHETI, AND M. NEIFELD, *Compressive imaging system design using task-specific information*, Applied Optics, 47 (2008), pp. 4457–4471.
- [2] P. BAHETI AND M. NEIFELD, *Recognition using information-optimal adaptive feature-specific imaging*, J. Opt. Soc. Amer. A, 26 (2009), pp. 1055–1070.
- [3] R.G. BARANIUK, V. CEVHER, M.F. DUARTE, AND C. HEGDE, *Model-based compressive sensing*, IEEE Trans. Inform. Theory, 56 (2010), pp. 1982–2001.
- [4] R. BARANIUK AND M. WAKIN, *Random projections of smooth manifolds*, Found. Comput. Math., 9 (2009), pp. 51–77.
- [5] S. BERGMAN AND B. OTTERSTEN, *Lattice-based linear precoding for MIMO channels with transmitter CSI*, IEEE Trans. Signal Process., (2008), pp. 2902–2914.
- [6] S. BOYD AND LIEVEN VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [7] E. CANDÈS, J. ROMBERG, AND T. TAO, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, IEEE Trans. Inform. Theory, 52 (2006), pp. 489–509.
- [8] E. CANDÈS AND M. WAKIN, *An introduction to compressive sampling*, vol. 25, IEEE Signal Process. Mag., (2008), pp. 21–30.
- [9] E. J. CANDÈS, J. K. ROMBERG, AND T. TAO, *Stable signal recovery from incomplete and inaccurate measurements*, Commun. Pure Appl. Math., 59 (2006), pp. 1207–1223.
- [10] W. R. CARSON, M. R. D. RODRIGUES, M. CHEN, L. CARIN, AND R. CALDERBANK, *How to focus the discriminative power of a dictionary*, in ICASSP '12, 2012.
- [11] M. CHEN, J. SILVA, J. PAISLEY, C. WANG, D. DUNSON, AND L. CARIN, *Compressive sensing on manifolds using a nonparametric mixture of factor analyzers: Algorithm and performance bounds*, vol. 58, IEEE Trans. Signal Process., (2010), pp. 6140–6155.
- [12] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, Wiley-Interscience, second ed., Aug. 1991.
- [13] D. L. DONOHO, *Compressed sensing*, IEEE Trans. Inform. Theory, 52 (2006), pp. 1289–1306.
- [14] J. DUARTE-CARVAJALINO, G. YU, L. CARIN, AND G. SAPIRO, *Adapted statistical compressive sensing: Learning to sense gaussian mixture models*, in ICASSP '12, Kyoto, Japan, 2012.
- [15] Y. ELDAR, P. KUPPINGER, AND H. BOLCSKEI, *Block-sparse signals: Uncertainty relations and efficient recovery*, IEEE Trans. Signal Process., 58 (2010), pp. 3042–3054.
- [16] D. ERDOGMUS AND J. C. PRINCIPE, *Lower and upper bounds for misclassification probability based on Renyi's information*, J. VLSI Signal Proc., 37 (2004), pp. 305–317.
- [17] M. E. GEHM, R. JOHN, D. J. BRADY, R. M. WILLETT, AND T. J. SCHULZ, *Single-shot compressive spectral imaging with a dual-disperser architecture*, Opt. Express, 15 (2007), pp. 14013–14027.
- [18] S. GOPARAJU, A. R. CALDERBANK, W. R. CARSON, M. R.D. RODRIGUES, AND F. PÉREZ-CRUZ, *When to add another dimension when communicating over MIMO channels*, in ICASSP '11, May 2011.
- [19] D. GUO, S. SHAMAI, AND S. VERDÚ, *Mutual information and minimum mean-square error in Gaussian channels*, IEEE Trans. Inform. Theory, (2005), pp. 1261–1282.
- [20] L. HE AND L. CARIN, *Exploiting structure in wavelet-based bayesian compressive sensing*, IEEE Trans. Signal Process., 57 (2009), pp. 3488–3497.
- [21] M. HELLMAN AND J. RAVIV, *Probability of error, equivocation, and the Chernoff bound*, IEEE Trans. Inform. Theory, 16 (1970), pp. 368–372.
- [22] Y. HITOMI, J. GU, M. GUPTA, T. MITSUNAGA, AND S. K. NAYAR, *Video from a single coded exposure photograph using a learned over-complete dictionary*, in IEEE International Conference on Computer Vision (ICCV), Nov 2011.
- [23] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, 2005.

- [24] S. JI, Y. XUE, AND L. CARIN, *Bayesian compressive sensing*, IEEE Trans. Signal Process., 56 (2008), pp. 2346–2356.
- [25] S. KAY, *Fundamentals of Statistical Signal Processing: Estimation Theory*, 2001.
- [26] CENK KÖSE AND RICHARD D. WESEL, *Universal Space-Time trellis codes*, IEEE Trans. Inform. Theory, 49 (2003), pp. 2717–2727.
- [27] M. LAMARCA, *Linear precoding for mutual information maximization in MIMO systems*, in Intl. Symposium on Wireless Communication Systems (ISWCS '09), Sept. 2009, pp. 26–30.
- [28] A. LOZANO, A. TULINO, AND S. VERDÚ, *Optimum power allocation for parallel Gaussian channels with arbitrary input distributions*, IEEE Trans. Inform. Theory, (2006), pp. 3033–3051.
- [29] M. LUSTIG, D. L. DONOHO, J. M. SANTOS, AND J. M. PAULY, *Compressed sensing MRI*, in IEEE Signal Process. Mag., 2007.
- [30] Z. NENADIC, *Information discriminant analysis: Feature extraction with an information-theoretic objective*, IEEE Trans. Pattern Anal. Mach. Intell., Jan. 2007, pp. 1394–1407.
- [31] D. PALOMAR AND Y. JIANG, *MIMO transceiver design via majorization theory*, in Foundations and Trends in Communications and Information Theory, Nov. 2006.
- [32] D. PALOMAR AND S. VERDÚ, *Gradient of mutual information in linear vector Gaussian channels*, IEEE Trans. Inform. Theory, Jan. 2006, pp. 141–154.
- [33] D. PÉREZ PALOMAR, J. CIOFFI, AND M. A. LAGUNAS, *Joint Tx-Rx beamforming design for multicarrier MIMO channels: A unified framework for convex optimization*, IEEE Trans. Signal Process., Sept. 2003, pp. 2381–2401.
- [34] M. PAYARÓ AND D. PÉREZ PALOMAR, *Hessian and concavity of mutual information, entropy, and entropy power in linear vector Gaussian channels*, IEEE Trans. Inform. Theory, Aug. 2009, pp. 3613–3628.
- [35] ———, *On optimal precoding in linear vector Gaussian channels with arbitrary input distribution*, in IEEE ISIT '09, July 2009, pp. 1085–1089.
- [36] F. PÉREZ-CRUZ, M. R.D. RODRIGUES, AND S. VERDÚ, *MIMO Gaussian channels with arbitrary inputs: Optimal precoding and power allocation*, IEEE Trans. Inform. Theory, Mar. 2010, pp. 1070–1084.
- [37] K. PETERSEN AND M. PEDERSEN, *The matrix cookbook*, in www2.imm.dtu.dk/pubdb/p.php?3274, 2008.
- [38] S. PRASAD, <http://arxiv.org/pdf/1010.1508v1.pdf>.
- [39] F. RENNA, *private communication*, 2012.
- [40] H. RICHTER, *Zur Abschätzung von Matrizennormen*, in Math. Nachr., 1958, pp. 178–187.
- [41] A. SCAGLIONE, P. STOICA, S. BARBAROSSA, G. B. GIANNAKIS, AND H. SAMPATH, *Optimal designs for space-time linear precoders and decoders*, IEEE Trans. Signal Process., 50 (2002), pp. 1051–1064.
- [42] L. SCARDOVI, *Information Based Control for State and Parameter Based Estimation*, PhD thesis, Genova University, <http://www.montefiore.ulg.ac.be/~scardovi/pub/phd.pdf>, Mar. 2005.
- [43] P. SCHNITER, *Exploiting structured sparsity in bayesian experimental design*, in CAMSAP '11, 2011.
- [44] M. SHANKAR, N.P. PITSIANIS, AND D.J. BRADY, *Compressive video sensors using multichannel imagers*, Appl. Opt., 49 (2010), pp. B9–B17.
- [45] T. SUN, C. LI, Y. ZHANG, L. XU, AND K. KELLY, *Compressive hyperspectral acquisition and endmember unmixing*, Proc. SPIE 8165, 81650D (2011).
- [46] B. TANG, J. TANG, AND Y. PENG, *MIMO radar waveform design in colored noise based on information theory*, IEEE Trans. Signal Process., Sept. 2010, pp. 4684–4697.
- [47] C. M. THEOBALD, *An inequality for the trace of the product of two symmetric matrices*, in Math. Proc. Camb. Phil. Soc., Oct. 1974, pp. 265–267.
- [48] H. S. WITSENHAUSEN, *A determinant maximization problem occurring in the theory of data communication*, in SIAM J. Appl. Math., Nov. 1975, pp. 515–522.
- [49] C. XIAO, Y. R. ZHENG, AND Z. DING, *Globally optimal linear precoders for finite alphabet signals over complex vector gaussian channels*, IEEE Trans. Signal Process., July 2011, pp. 3301–3314.

Appendix A. Complex Derivatives and Gradients.

Throughout the paper we adopt the definition of the formal partial complex derivative of a real-valued scalar function f with respect to a complex-valued variable \mathbf{x} given by [32] [49]:

$$\frac{\partial f}{\partial x^*} \triangleq \frac{1}{2} \left[\frac{\partial f}{\partial \text{Re}(x)} + j \frac{\partial f}{\partial \text{Im}(x)} \right] \quad (\text{A.1})$$

The definition of the complex gradient of a real-valued function f with respect to a complex-valued matrix \mathbf{X} is given by:

$$\nabla_{\mathbf{X}} f \triangleq \frac{\partial f}{\partial \mathbf{X}^*} \quad (\text{A.2})$$

where $[\nabla_{\mathbf{X}} f]_{ij} = \partial f / \partial [\mathbf{X}^*]_{ij}$.

Appendix B. Helpful Lemmas.

In the proofs of the Theorems stated in this paper we will find the following lemmas helpful

LEMMA B.1 (Sylvester's Determinant Theorem). *We have a "cyclic" property of determinants for two matrices $\mathbf{A} \in \mathcal{C}^{n \times m}$ and $\mathbf{B} \in \mathcal{C}^{m \times n}$:*

$$\det(\mathbf{I}_n + \mathbf{A}\mathbf{B}) = \det(\mathbf{I}_m + \mathbf{B}\mathbf{A}). \quad (\text{B.1})$$

In the following four lemmas we denote two Hermitian matrices by $\Sigma_{\mathbf{A}}, \Sigma_{\mathbf{B}} \in \mathcal{C}^{m \times m}$ which have eigenvalues $\lambda_1 \geq \dots \geq \lambda_m$ and $\mu_1 \geq \dots \geq \mu_m$, respectively. The eigenvalue decomposition of these two matrices are $\Sigma_{\mathbf{A}} = \mathbf{U}_{\mathbf{A}} \Lambda_{\mathbf{A}} \mathbf{U}_{\mathbf{A}}^\dagger$ and $\Sigma_{\mathbf{B}} = \mathbf{U}_{\mathbf{B}} \Lambda_{\mathbf{B}} \mathbf{U}_{\mathbf{B}}^\dagger$, where $\mathbf{U}_{\mathbf{A}}, \mathbf{U}_{\mathbf{B}} \in \mathbb{S}^{m \times m}$, $\Lambda_{\mathbf{A}} = \text{Diag}(\lambda_1, \dots, \lambda_m)$ and $\Lambda_{\mathbf{B}} = \text{Diag}(\mu_1, \dots, \mu_m)$.

LEMMA B.2 (Theorem 1.3.12 in [23]). *The matrices $\Sigma_{\mathbf{A}}$ and $\Sigma_{\mathbf{B}}$ commute if and only if they are simultaneously diagonalizable, i.e., both $\mathbf{U}\Sigma_{\mathbf{A}}\mathbf{U}^\dagger$ and $\mathbf{U}\Sigma_{\mathbf{B}}\mathbf{U}^\dagger$ are diagonal matrices for some unitary matrix \mathbf{U} .*

LEMMA B.3 (Richter [40]).

$$\sum_{i=1}^m \lambda_i \mu_{m+1-i} \leq \text{tr}(\Sigma_{\mathbf{A}} \Sigma_{\mathbf{B}}) \leq \sum_{i=1}^m \lambda_i \mu_i. \quad (\text{B.2})$$

REMARK 2. *Sufficient conditions for achieving the upper and lower bounds are $\mathbf{U}_{\mathbf{A}} = \mathbf{U}_{\mathbf{B}}$ and $\mathbf{U}_{\mathbf{A}} = \mathbf{U}_{\mathbf{B}} \mathbf{J}_m$, respectively. The sufficient condition to achieve the lower bound was given by Köse and Wesel in Theorem 2 in [26] and Theobald [47] also gave necessary and sufficient conditions for achieving the upper bound, which allow for the multiplicity of eigenvalues.*

LEMMA B.4 (Lemma 3 in Witzenhausen [48]).

$$\det(\mathbf{I}_m + \Sigma_{\mathbf{A}} \Sigma_{\mathbf{B}}) \leq \prod_{i=1}^m (1 + \lambda_i \mu_i). \quad (\text{B.3})$$

REMARK 3. A sufficient condition for achieving the upper bound is $\mathbf{U}_{\mathbf{A}} = \mathbf{U}_{\mathbf{B}}$. Witzenhausen gave further sufficient conditions which allow for the multiplicity of eigenvalues, stating that if equality holds then $\Sigma_{\mathbf{A}}$ and $\Sigma_{\mathbf{B}}$ commute and the diagonalizing matrix is such that the eigenvalues are aligned in the same order.

LEMMA B.5. Let $\mathbf{P} \in \mathbb{C}^{m \times n}$ denote a rectangular matrix, $\Sigma_{\mathbf{H}} \in \mathbb{C}^{n \times n}$ denote a positive semi-definite matrix, and $\mathbf{P}^\dagger \Sigma_{\mathbf{H}} \mathbf{P} \in \mathbb{C}^{m \times m}$ be a diagonal matrix with diagonal elements in decreasing order (possibly with some zero diagonal elements). Then, there is a matrix of the form $\bar{\mathbf{P}} = \mathbf{V}_{\mathbf{H}} [\mathbf{\Lambda}, \mathbf{0}]$ that satisfies:

$$\bar{\mathbf{P}}^\dagger \Sigma_{\mathbf{H}} \bar{\mathbf{P}} = \alpha \mathbf{P}^\dagger \Sigma_{\mathbf{H}} \mathbf{P} \quad (\text{B.4})$$

$$\text{tr}(\bar{\mathbf{P}} \bar{\mathbf{P}}^\dagger) = \text{tr}(\mathbf{P} \mathbf{P}^\dagger) \quad (\text{B.5})$$

where $\alpha \geq 1$, $\mathbf{V}_{\mathbf{H}}$ is a unitary matrix with columns equal to the eigenvectors of matrix $\Sigma_{\mathbf{H}}$ corresponding to the $\min(n, m)$ largest eigenvalues in decreasing order and $\mathbf{\Lambda}$ is square diagonal matrix of size $\min(n, m)$.

Proof. This is a modification of Lemma 3.16 in [31]. \square

LEMMA B.6. For the complex gradient defined in (A.2) and general matrices $\mathbf{A} \in \mathbb{C}^{m \times m}$, $\mathbf{B} \in \mathbb{C}^{n \times n}$ and $\mathbf{X} \in \mathbb{C}^{m \times n}$, we have:

$$\nabla_{\mathbf{X}} \text{tr}(\mathbf{A} \mathbf{X} \mathbf{B} \mathbf{X}^\dagger) = \mathbf{A} \mathbf{X} \mathbf{B} \quad (\text{B.6})$$

Proof. Using properties of differentials (32) and (33) from [37] we have:

$$\partial \text{tr}(\mathbf{A} \mathbf{X} \mathbf{B} \mathbf{X}^\dagger) = \text{tr}[\mathbf{A} \partial(\mathbf{X}) \mathbf{B} \mathbf{X}^\dagger] + \text{tr}[\mathbf{A} \mathbf{X} \mathbf{B} \partial(\mathbf{X}^\dagger)]. \quad (\text{B.7})$$

Together with the results for complex derivatives (219), (220), (221) and (222) from [37] we have:

$$\frac{\partial}{\partial \text{Re}(\mathbf{X})} \text{tr}(\mathbf{A} \mathbf{X} \mathbf{B} \mathbf{X}^\dagger) = \mathbf{A}^\top \mathbf{X}^* \mathbf{B}^\top + \mathbf{A} \mathbf{X} \mathbf{B} \quad (\text{B.8})$$

$$i \frac{\partial}{\partial \text{Im}(\mathbf{X})} \text{tr}(\mathbf{A} \mathbf{X} \mathbf{B} \mathbf{X}^\dagger) = -\mathbf{A}^\top \mathbf{X}^* \mathbf{B}^\top + \mathbf{A} \mathbf{X} \mathbf{B} \quad (\text{B.9})$$

and the result follows. \square

REMARK 4. This is the counterpart for complex-valued matrices to result (108) in [37] for real-valued matrices; note that the term $\mathbf{A}^\top \mathbf{X} \mathbf{B}^\top$ is absent in the complex case.

LEMMA B.7. For the complex gradient defined in (A.2) and general matrices $\mathbf{A} \in \mathbb{C}^{m \times m}$, $\mathbf{B} \in \mathbb{C}^{n \times n}$, $\mathbf{C} \in \mathbb{C}^{n \times n}$ and $\mathbf{X} \in \mathbb{C}^{m \times n}$, we have:

$$\begin{aligned} \nabla_{\mathbf{X}} \text{tr}[(\mathbf{A} + \mathbf{X} \mathbf{B} \mathbf{X}^\dagger)^{-1} (\mathbf{X} \mathbf{C} \mathbf{X}^\dagger)] \\ = (\mathbf{A} + \mathbf{X} \mathbf{B} \mathbf{X}^\dagger)^{-1} \mathbf{X} \mathbf{C} [\mathbf{I} - \mathbf{X}^\dagger (\mathbf{A} + \mathbf{X} \mathbf{B} \mathbf{X}^\dagger)^{-1} \mathbf{X} \mathbf{B}] \end{aligned} \quad (\text{B.10})$$

Proof. Using properties of differentials (32), (33) and (36) in [37] and the abbreviation $\mathbf{Y} = \mathbf{A} + \mathbf{X} \mathbf{B} \mathbf{X}^\dagger$, we have:

$$\begin{aligned} \partial \text{tr} [\mathbf{Y}^{-1} (\mathbf{X} \mathbf{C} \mathbf{X}^\dagger)] &= \text{tr} \{ \mathbf{Y}^{-1} \partial (\mathbf{X} \mathbf{C} \mathbf{X}^\dagger) \} \\ &+ \text{tr} \{ -\mathbf{Y}^{-1} \partial (\mathbf{A} + \mathbf{X} \mathbf{B} \mathbf{X}^\dagger) \mathbf{Y}^{-1} (\mathbf{X} \mathbf{C} \mathbf{X}^\dagger) \} \end{aligned}$$

Applying Lemma B.6 the result follows. \square

REMARK 5. *This is the counterpart for complex-valued matrices to result (116) in [37] for real-valued matrices; note that we do not require the assumption that \mathbf{B} and \mathbf{C} are Hermitian (symmetric) and there is no factor of 2.*

Appendix C. Proof of Theorem 4.1.

Proof. We first provided an alternative derivation of this proof in [10]. The current proof is derived directly from the proof for an identical theorem for radar in [46]. We restate the mutual information between the input and output of the compressive sensing model in (3.1) for a multivariate Gaussian source as follows:

$$\mathcal{I}(\mathbf{x}; \mathbf{y}) = \log \det (\mathbf{I}_m + \mathbf{M}^\dagger \boldsymbol{\Sigma}_{\mathbf{w}}^{-1} \mathbf{M} \boldsymbol{\Sigma}_{\mathbf{x}}). \quad (4.2)$$

Note that for a unitary matrix \mathbf{U} , the kernel $\mathbf{P} = \mathbf{M}\mathbf{U}$ has the same power as \mathbf{M} , i.e., $\text{tr}(\mathbf{P}\mathbf{P}^\dagger) = \text{tr}(\mathbf{M}\mathbf{M}^\dagger)$, but it may have different mutual information. In particular, a choice of \mathbf{U} that maximizes the mutual information for a given \mathbf{M} is $\mathbf{U} = \mathbf{U}_{\mathbf{x}}$. This can be seen from Lemma B.4 and Remark B.4. From Lemma B.5 we know that there exists a matrix $\bar{\mathbf{P}} = \mathbf{U}_{\mathbf{w}} [\mathbf{A}, \mathbf{0}]$, which satisfies $\text{tr}(\bar{\mathbf{P}} \bar{\mathbf{P}}^\dagger) = \text{tr}(\mathbf{P}\mathbf{P}^\dagger)$ and $\bar{\mathbf{P}}^\dagger \boldsymbol{\Sigma}_{\mathbf{H}} \bar{\mathbf{P}} = \alpha \mathbf{P}^\dagger \boldsymbol{\Sigma}_{\mathbf{w}}^{-1} \mathbf{P}$ where $\alpha \geq 1$. Since the function $\det(\mathbf{I} + \alpha \mathbf{A})$ is monotonically increasing in α for a positive semi-definite matrix \mathbf{A} , the optimal kernel matrix must have the form of $\mathbf{M}^* = \mathbf{U}_{\mathbf{M}}^* \mathbf{A}_{\mathbf{M}}^* \mathbf{V}_{\mathbf{M}}^{*\dagger} = \mathbf{U}_{\mathbf{w}} [\mathbf{A}, \mathbf{0}] \mathbf{U}_{\mathbf{x}}^\dagger$.

Finally, we determine the optimal singular values by optimizing the mutual information with respect to the eigenvalues rather than the singular values, since 1) they map one-to-one (up to a factor of $\exp j\theta$, which does not affect the mutual information) and 2) this new optimization problem is convex, so the Karush-Kuhn-Tucker (KKT) optimality conditions [6] define the unique global optimum. This is given by:

$$\lambda_{M_i}^* = \begin{cases} 0, & \frac{1}{\eta} - \frac{\lambda_{w_i}}{\lambda_{x_i}} \leq 0, \\ \frac{1}{\eta} - \frac{\lambda_{w_i}}{\lambda_{x_i}}, & \frac{1}{\eta} - \frac{\lambda_{w_i}}{\lambda_{x_i}} > 0 \end{cases} \quad (C.1)$$

where η is such that the average unit norm row constraint is satisfied, i.e., $\frac{1}{\ell} \sum \lambda_{M_i}^* = 1$, where the eigenvalues of $\boldsymbol{\Sigma}_{\mathbf{x}}$ are arranged in descending order and the eigenvalues of $\boldsymbol{\Sigma}_{\mathbf{w}}$ are arranged in ascending order, i.e., $\lambda_{x_1} \geq \dots \geq \lambda_{x_m} \geq 0$ and $0 \leq \lambda_{w_1} \leq \dots \leq \lambda_{w_\ell}$. \square

Appendix D. Proof of Theorem 4.2.

Proof. The proof draws on the work by Payaró and Palomar [35], which described the generalized mercury waterfilling aspect of the solution, but not the mode alignment aspect, and the work by Lamarca [27], which described the mode alignment aspect but did not focus on the generalized mercury waterfilling interpretation. The current proof highlights both the mode alignment and mercury waterfilling aspects of the solution.

The solution to the optimization problem in (4.1) satisfies the KKT optimality conditions:

$$\nabla_{\mathbf{M}} \left\{ -\mathcal{I}(\mathbf{x}; \mathbf{y}) - \eta \cdot \left[\ell - \text{tr}(\mathbf{M}\mathbf{M}^\dagger) \right] \right\} \Big|_{\mathbf{M}=\mathbf{M}^*} = 0 \quad (\text{D.1})$$

$$\eta \cdot \left[\ell - \text{tr}(\mathbf{M}^* \mathbf{M}^{*\dagger}) \right] = 0 \quad (\text{D.2})$$

with $\eta \geq 0$. Using the relationship between the gradient of the mutual information and the MMSE matrix in (3.7), the optimal kernel satisfies:

$$\eta \cdot \mathbf{M}^* \mathbf{M}^{*\dagger} = \Sigma_{\mathbf{w}}^{-1} (\mathbf{M}^* \mathbf{E}^* \mathbf{M}^{*\dagger}). \quad (\text{D.3})$$

We note that (D.3) is diagonalized by $\mathbf{U}_{\mathbf{M}}^*$, by definition, from which it can be seen that the matrices $\Sigma_{\mathbf{w}}^{-1}$ and $\mathbf{M}^* \mathbf{E}^* \mathbf{M}^{*\dagger}$ commute. From this observation, together with the fact that (D.3) is Hermitian and Lemma B.2, we deduce that:

$$\mathbf{U}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{w}} \mathbf{\Pi}_{\mathbf{U}}^* \mathbf{\Lambda}_{\mathbf{U}} \quad (\text{D.4})$$

where $\mathbf{\Lambda}_{\mathbf{U}}$ is a diagonal matrix with unit modulus diagonal elements and $\mathbf{\Pi}_{\mathbf{U}}^*$ is a permutation matrix. Furthermore, $\mathbf{U}_{\mathbf{M}}^{*\dagger} \mathbf{M}^* \mathbf{E}^* \mathbf{M}^{*\dagger} \mathbf{U}_{\mathbf{M}}^*$ is a diagonal matrix, from which we can infer:

$$\mathbf{V}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}_{\mathbf{V}}^* \mathbf{\Lambda}_{\mathbf{V}} \quad (\text{D.5})$$

where $\mathbf{\Lambda}_{\mathbf{V}}$ is a diagonal matrix with unit modulus diagonal elements and $\mathbf{\Pi}_{\mathbf{V}}^*$ is a permutation matrix. Both mutual information and the MMSE matrix are independent of $\mathbf{\Lambda}_{\mathbf{U}}$ and $\mathbf{\Lambda}_{\mathbf{V}}$, allowing us to write without loss of generality the optimal unitary matrices as follows:

$$\mathbf{U}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{w}} \quad (\text{D.6})$$

$$\mathbf{V}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}^* \quad (\text{D.7})$$

where $\mathbf{\Pi}^*$ is some optimal permutation matrix.

By setting $\mathbf{U}_{\mathbf{M}}^* = \mathbf{U}_{\mathbf{w}}$ we can now obtain an equivalent⁸ channel model:

$$\mathbf{y}' = \mathbf{\Lambda}_{\mathbf{w}}^{-1/2} \mathbf{\Lambda}_{\mathbf{M}} \mathbf{V}_{\mathbf{M}}^\dagger \mathbf{x} + \mathbf{n} \quad (\text{D.8})$$

where $\mathbf{y}' = \mathbf{\Lambda}_{\mathbf{w}}^{-1/2} \mathbf{U}_{\mathbf{w}}^\dagger \mathbf{y}$ and $\mathbf{n} = \mathbf{\Lambda}_{\mathbf{w}}^{-1/2} \mathbf{U}_{\mathbf{w}}^\dagger \mathbf{w}$ is zero-mean circularly symmetric complex Gaussian noise with identity covariance, $\mathbf{n} \sim \mathcal{CN}(\mathbf{n}; \mathbf{0}, \mathbf{I})$.

It was shown in [49] that for a fixed value of $\mathbf{V}_{\mathbf{M}}$ the mutual information $\mathcal{I}(\mathbf{x}; \mathbf{y}')$ is concave with respect to the squared singular values of $\mathbf{\Lambda}_{\mathbf{M}}$, i.e., the following optimization problem has a unique global optimum given by the KKT conditions, where $\mathbf{\Lambda}_{\mathbf{Q}} = \mathbf{\Lambda}_{\mathbf{M}} \mathbf{\Lambda}_{\mathbf{M}}^\dagger$:

$$\begin{aligned} & \underset{\lambda_{M_1}, \lambda_{M_2}, \dots, \lambda_{M_\ell}}{\text{maximize}} && \mathcal{I}(\mathbf{x}; \mathbf{y}') \\ & \text{subject to} && \sum_{i=1}^{\ell} \lambda_{M_i} \leq \ell \\ & && \lambda_{M_i} \geq 0 \end{aligned} \quad (\text{D.9})$$

⁸The equivalence is in the sense that the mutual information between the input and the output of both models is equal, i.e., $\mathcal{I}(\mathbf{x}; \mathbf{y}) = \mathcal{I}(\mathbf{x}; \mathbf{y}')$.

The Lagrangian for this optimization problem is:

$$\mathcal{L}(\Lambda_{\mathbf{Q}}) = \mathcal{I}(\mathbf{x}; \mathbf{y}') + \eta \left(\ell - \sum_{i=1}^{\ell} \lambda_{M_i} \right) + \sum_{i=1}^{\ell} \eta_i \lambda_{M_i} \quad (\text{D.10})$$

and the Karush-Kuhn-Tucker conditions state that:

$$\left. \frac{\partial}{\partial \Lambda_{\mathbf{Q}}} \mathcal{L}(\Lambda_{\mathbf{Q}}) \right|_{\Lambda_{\mathbf{Q}} = \Lambda_{\mathbf{Q}}^*} = 0 \quad (\text{D.11})$$

$$\eta \cdot \left(\ell - \sum_{i=1}^{\ell} \lambda_{M_i}^* \right) = 0 \quad (\text{D.12})$$

$$\eta_i \cdot \lambda_{M_i}^* = 0, \quad i = 1, \dots, \ell \quad (\text{D.13})$$

By using the result from [49] that states that:

$$\frac{\partial}{\partial \Lambda_{\mathbf{Q}}} \mathcal{I}(\mathbf{x}; \mathbf{y}') = \text{Diag}(\mathbf{E}_{\mathbf{x}'} \Lambda_{\mathbf{w}}^{-1}) \quad (\text{D.14})$$

where $\mathbf{E}' = \mathbf{V}_{\mathbf{M}}^{\dagger} \mathbf{E} \mathbf{V}_{\mathbf{M}}$ is the MMSE matrix associated with the estimation of $\mathbf{x}' = \mathbf{V}_{\mathbf{M}}^{\dagger} \mathbf{x}$ from \mathbf{y}' , it is possible to rewrite (D.11) as follows:

$$\eta \lambda_{w_i} - \text{mmse}_i(\mathbf{V}_{\mathbf{M}}, \Lambda_{\mathbf{Q}}^*) = \eta_i \lambda_{w_i}, \quad i = 1, \dots, \ell. \quad (\text{D.15})$$

where $\text{mmse}_i(\mathbf{V}_{\mathbf{M}}, \Lambda_{\mathbf{Q}})$ denotes the i -th diagonal entry of $\mathbf{E}_{\mathbf{x}'}$ for that particular $\Lambda_{\mathbf{Q}}$.

From the KKT conditions, we know that if $\lambda_{M_i}^* > 0$ then $\eta_i = 0$ and that $\eta > 0$. For a given value of η , the value of $\lambda_{M_i}^*$ can be calculated from the relationship $\eta \lambda_{w_i} = \text{mmse}_i(\mathbf{V}_{\mathbf{M}}, \Lambda_{\mathbf{Q}}^*)$ for fixed $\lambda_{M_j}, \forall j \neq i$. The function mmse_i is non-negative and monotonically decreasing in $\lambda_{M_i} \in [0, \infty]$ for fixed $\lambda_{M_j}, \forall j \neq i$, and its maximum value is given when $\lambda_{M_i} = 0$. Therefore if $\eta \lambda_{w_i} > \text{mmse}_i(\mathbf{V}_{\mathbf{M}}, \Lambda_{\mathbf{Q}}^*|_{\lambda_{M_i}^*=0})$ where $\Lambda_{\mathbf{Q}}^*|_{\lambda_{M_i}^*=0} = \text{Diag}(\lambda_{M_1}^*, \dots, \lambda_{M_{i-1}}^*, 0, \lambda_{M_{i+1}}^*, \dots, \lambda_{M_{\ell}}^*)$, then $\lambda_{M_i}^* = 0$ and $\eta_i \neq 0$.

This result is true for all values of $\mathbf{V}_{\mathbf{M}}$, therefore it is also true when $\mathbf{V}_{\mathbf{M}} = \mathbf{U}_{\mathbf{E}}^* \mathbf{\Pi}^*$ and the result follows. \square

Appendix E. Proof of Theorem 5.1.

Proof. Note that:

$$p(\mathbf{y}) = \mathcal{CN}(\mathbf{y}; \mathbf{M} \mathbf{x}, \Sigma_{\mathbf{w}} + \mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^{\dagger}) \quad (\text{E.1})$$

and so:

$$p^{\alpha}(\mathbf{y}) = \frac{\mathcal{CN}(\mathbf{y}; \mathbf{M} \mathbf{x}, \frac{1}{\alpha}(\Sigma_{\mathbf{w}} + \mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^{\dagger}))}{\alpha^k (2\pi)^{k(\alpha-1)} \det(\Sigma_{\mathbf{w}} + \mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^{\dagger})^{\alpha-1}} \quad (\text{E.2})$$

By substituting this into the expression for Rényi entropy it follows that:

$$h_{\alpha}(\mathbf{y}) = \log \left[(2\pi)^k \det(\Sigma_{\mathbf{w}} + \mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^{\dagger}) \right] + \frac{k \log \alpha}{(\alpha - 1)}. \quad (\text{E.3})$$

The result now follows using the definition of Shannon entropy for Gaussian sources. \square

Appendix F. Proofs for Gradient of Rényi entropy.

Proof. Let us first show that we can express the following relevant gradient analytically:

$$\begin{aligned} \nabla_{\mathbf{M}} \log \mathcal{CN}(\mathbf{0}; \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j}) = \\ - \nabla_{\mathbf{M}} \left(k \log 2\pi \right) - \nabla_{\mathbf{M}} \{ \log \det \boldsymbol{\Sigma}_{i,j} \} \end{aligned} \quad (\text{F.1})$$

$$- \nabla_{\mathbf{M}} \{ \boldsymbol{\mu}_{i,j}^{\top} \boldsymbol{\Sigma}_{i,j}^{-1} \boldsymbol{\mu}_{i,j} \}, \quad (\text{F.2})$$

where $\boldsymbol{\mu}_{i,j} = \mathbf{M}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)$ and $\boldsymbol{\Sigma}_{i,j} = \mathbf{M}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) \mathbf{M}^{\dagger} + 2\boldsymbol{\Sigma}_{\mathbf{w}}$.

The first term is zero, the second term is the mutual information for a complex Gaussian distribution and can be evaluated using (3.7), relating the mutual information and MMSE matrix:

$$\nabla_{\mathbf{M}} \log \det \boldsymbol{\Sigma}_{i,j} = (2\boldsymbol{\Sigma}_{\mathbf{w}})^{-1} \mathbf{M} \mathbf{E}_{i,j} \quad (\text{F.3})$$

where the $\mathbf{E}_{i,j} = \left[(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j)^{-1} + \mathbf{M}^{\dagger} (2\boldsymbol{\Sigma}_{\mathbf{w}})^{-1} \mathbf{M} \right]^{-1}$ is the MMSE matrix if the input signal \mathbf{x} was Gaussian distributed with covariance $(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j)$ and distorted by Gaussian noise with covariance $2\boldsymbol{\Sigma}_{\mathbf{w}}$. It can also be expressed:

$$\nabla_{\mathbf{M}} \log \det \boldsymbol{\Sigma}_{i,j} = \boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M} (\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) \quad (\text{F.4})$$

where we can use Woodbury's Inversion Lemma to convert between the two. The third and final term, using Lemma B.7 and chain rule (94) in [32], can be expressed:

$$\begin{aligned} \nabla_{\mathbf{M}} \{ \boldsymbol{\mu}_{i,j}^{\top} \boldsymbol{\Sigma}_{i,j}^{-1} \boldsymbol{\mu}_{i,j} \} = \boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M} (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^{\dagger} \\ \times \{ \mathbf{I} - \mathbf{M}^{\dagger} \boldsymbol{\Sigma}_{i,j}^{-1} \mathbf{M} (\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) \}. \end{aligned} \quad (\text{F.5})$$

□